

OPERATING CFDP IN THE INTERPLANETARY INTERNET

Scott C. Burleigh

Jet Propulsion Laboratory

California Institute of Technology, Pasadena, CA

Email: Scott.Burleigh@jpl.nasa.gov

Abstract

CFDP (CCSDS File Delivery Protocol) is a new international standard developed to meet a comprehensive set of deep space file transfer requirements, as articulated by a number of space agencies including NASA, ESA, NASDA, CNES, and BNSC/DERA. In addition, CFDP will in effect serve as a prototype for the future Interplanetary Internet (IPN) as envisioned by the IPN Study team: it encompasses a subset of the anticipated functionality of the IPN, and it implements several key IPN design concepts including store-and-forward operation with deferred transmission and concurrent transactions. The IPN design goes beyond the capabilities of CFDP in some key areas, notably network scalability and compatibility with anticipated terrestrial delay-tolerant networking technology, but deployed IPN protocols will complement rather than supersede CFDP. This paper examines the design elements of these technologies that will simplify their integration and discusses the resulting new capabilities, such as efficient transmission of large files via multiple relay satellites operating in parallel.¹

Introduction

CFDP is the CCSDS File Delivery Protocol [1], an international standard for automatic, reliable file transfer between spacecraft and ground (in both directions), built on the familiar Consultative Committee for Space Data Systems (CCSDS) [2] space data communication protocols. It is designed to support the operation of spacecraft by means of file transfer and remote file system management, in satisfaction of requirements developed by subpanel 1F of CCSDS. Its capabilities include:

- Reliably copying a file from the filestore of one *entity* (protocol engine, located in a spacecraft or ground control center) to that of another entity.
- Reliable transmission of arbitrary small messages, defined by the user, in the *metadata* accompanying a file.
- Reliable transmission of file system management commands to be executed automatically at a remote entity – typically at a spacecraft – upon complete reception of a file.

CFDP is designed to offer these capabilities even across interplanetary distances. As such, it must function despite extremely long data propagation delays (measured in minutes or hours, rather than in milliseconds as in terrestrial networks) and frequent, lengthy interruptions in connectivity.

Design Concepts

It is expected that a pair of CFDP entities which have files to exchange may at any given moment be unable to communicate; for example, a spacecraft orbiting Mars may be on the far side of the planet, unable to transmit to Earth. For this reason, CFDP is built entirely on a *store and forward* communication model. If transmission of a file from Earth to a Mars-orbiting spacecraft is interrupted when the spacecraft passes behind the planet, the CFDP entities at both ends of the transmission simply

¹ This work was performed at and under the direction of the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

store their outbound protocol data units (*PDU*s) – possibly in non-volatile memory, to assure continued service even in the event of an unplanned system reset – until the spacecraft re-emerges and transmission can resume. A collateral benefit of this model is that it largely insulates user applications from the state of the communication system: an instrument can record an observation in a file and “transmit” it (that is, submit it to CFDP for transmission) to Earth immediately without considering whether or not physical transmission is currently possible. By sequestering outbound data management and transmission planning functions within CFDP, this *deferred transmission* can simplify flight and ground software and thereby reduce mission costs.

Using powerful forward error correction coding can minimize data loss in communication across deep space but cannot eliminate it altogether. Consequently CFDP supports optional “acknowledged” modes of operation in which data loss is automatically detected and retransmission of the lost data is automatically requested. However, the large signal propagation delays that characterize interplanetary transmission limit the usefulness of the retransmission strategies commonly used in terrestrial protocols. For example, delaying the transmission of *PDU N* until an acknowledgment that *PDU N - 1* – or even *PDU N - 100* – has been received would significantly retard data flow if the round trip time on the link exceeded the time required to radiate the *PDU*(s) for which acknowledgment is required. For this reason, CFDP’s retransmission model is one of *concurrent transmission*: data *PDU*s for multiple files are transmitted as rapidly as possible, one after another, without waiting for acknowledgment, and requests for retransmission are handled asynchronously as they are received. As a result, portions of multiple files may be in transit concurrently.

Implied in concurrent transmission is the possibility that data may arrive out of the order in which they were sent: *PDU*s that were lost in initial transmission and subsequently retransmitted will typically arrive long after *PDU*s that were initially transmitted at a later time but were received without incident. Because data arrival order can’t reliably be used to reconstruct files, all CFDP *PDU*s are tagged with *transaction* identifiers that link them to the files of which they are constituents. Moreover, because in some cases a part of a file (such as a camera image) may be useful even before any missing portions of the file arrive, CFDP offers optional *out-of-order delivery* – that is, delivery of data to the user in the order in which it arrives rather than the order in which it was transmitted.

For simplicity, the facilities of any single CFDP entity are offered directly only to a single *user application*. Variability in mission operations policy can normally be confined to the design and implementation of the user application, so that the CFDP implementation may be re-used without modification in any number of mission contexts. The user application itself is, of course, free to multiplex CFDP service access indirectly to any number of other applications. In addition, the CFDP Recommendation defines several standard capabilities that may be implemented in a given user application to enhance CFDP without modifying it; these standard “user operations” can be thought of as additional protocols layered on top of CFDP.

CFDP is designed for re-use without modification in any number of communication environments as well. No specific direct interface to radio hardware, or even to any specific link-layer protocol, is mandated in the Recommendation. Instead, an abstract underlying “unitdata transfer” or *UT layer* service is assumed to be available for CFDP’s use. The specific requirements imposed on the *UT layer* are minimal, enabling CFDP to be run on top of such widely varying services as UDP/IP on the Internet (the configuration chosen for interoperability testing of different space agencies’ implementations of CFDP) and the Packet services of CCSDS Telemetry [3] and Telecommand [4] in space. Each CFDP *PDU* is simply encapsulated in a single *UT* “service data unit” (e.g., packet or datagram) for transmission, and the details of conveying the *PDU*s from one CFDP entity to another are left to the *UT* service.

Other Operational Features

The CFDP Recommendation is largely silent on priority assignment, bandwidth sharing, and other details of transmission ordering. *Flow labels* can be associated with transactions, but interpreting and responding to flow labels is left to the implementation.

The files that CFDP transmits may be simple arrays of octets or they may be structured as arrays of variable-length *records*. No particular record format is dictated by the CFDP Recommendation. However, if the user application indicates that a given file which is to be transmitted is a record-structured file, the CFDP entity is obliged to assure that the first octet of each file data segment PDU is the first octet of a record; the manner in which record boundaries are detected is left to the implementation.

For simplicity, CFDP offers only a basic “put file” service. However, standardized user operations built on this basic service can be used to:

- “get file” from a remote entity
- instruct a remote entity to send a file to some other remote entity
- list the contents of a remote entity’s filestore
- query the status of a transaction at a remote entity

CFDP Basic Deployment in Core Architecture

In the simplest operational scenario for CFDP, a mission control center on Earth exchanges files directly with a single spacecraft. Although the CFDP communication in this case is “end to end” because the spacecraft and mission control center are the original source and final destination of every file, it is also “point to point” in the sense that it operates over what is in concept a single communications link.

Nominally the actual link in this topology is a single physical link, provided by mutual radio transmission between the two entities, but the same scenario applies when the two entities are (for example) communicating via a UT layer service that uses UDP over the Internet: the multiple physical links bridged by the underlying IP protocol are invisible and functionally irrelevant to the CFDP entities.

Operation in this scenario is possible only when a number of key constraints are met:

- The link must be topologically possible. In deep space this means that the communicating endpoint entities must be in line of sight of each other.
- The link must be operationally possible. In deep space this means that the spacecraft must have been commanded to point its antenna at Earth and turn on its transponder at a specific moment and for a specific length of time, and analogous pointing and system configuration must have been commanded on the ground for the same time frame.
- The link must be electronically possible. In deep space this means that the sending entity must have sufficient transmitter power to get a signal to the receiver and the receiving entity must have sufficient receiver power to acquire that signal.

An example of operation in this configuration is the Deep Impact mission [5] scheduled for launch in 2004. Deep Impact will send a pair of spacecraft to investigate the comet Tempel 1. One of the spacecraft, termed the “impactor”, will crash into the comet; a nearby “flyby” craft will photograph and analyze the material ejected due to the impact. CFDP will be used to operate the flyby craft from Earth: command sequence files will be sent in acknowledged mode from Earth to the spacecraft, and files of science observations will be sent from the spacecraft to Earth in unacknowledged mode.

For operation in this basic configuration, the core elements of CFDP are sufficient: simple user applications at both ends of the link, a single UT layer implementation built on the CCSDS Telemetry and Telecommand link-layer protocols, and CFDP's file system functions and point-to-point retransmission capability.

CFDP Advanced Deployment in Extended Architecture

The store-and-forward model on which CFDP is built will enable it to be operated in more complex scenarios as well. CFDP is especially well suited to operations involving relay satellites in deep space. For example, a Mars rover might have insufficient transmitter power to send files directly to Earth at high speed yet could easily transfer the files to a relay satellite in Mars orbit; that satellite, equipped with larger solar power panels and a more powerful radio, could then forward the files to Earth. The transmission from the rover to the satellite might occur at a time when both are on the far side of the planet, out of sight of Earth, but this would present no problem since CFDP on board the satellite will simply store the files until the spacecraft emerges from Mars' shadow.

More complex relaying scenarios are equally plausible: there might be multiple rovers and landers, multiple relay satellites, and perhaps even smaller devices that would relay files via rovers because they lack enough power even to transmit to the satellites.

All such scenarios are rooted in the violation of one or more of the constraints on basic deployment described above, forcing each end-to-end transmission to entail a series of two or more point-to-point exchanges of data. For this purpose, the core CFDP architecture must be extended. Proposed revisions to the CFDP Recommendation that are currently under review offer two alternative solutions:

- A *store-and-forward overlay* system may be added to user applications as a new standard user operation. The user application at each relay point examines each incoming file and, if the accompanying metadata indicates that the file's final destination is elsewhere, initiates another point-to-point file transmission either to the final destination or to another relay point that is farther along the route.
- *Extended procedures* may be added to CFDP itself. The CFDP entity at each relay point checks the final destination of each incoming file and, if necessary, initiates another point-to-point transmission toward that destination; the file is never delivered to the user application at any relay point.

Multiple elements of CFDP's architecture may in fact require extension for these scenarios. Complexity must be added either to the user applications (store-and-forward overlay) or to CFDP (the extended procedures). In addition, multiple UT layer implementations may be needed: for example, CFDP communication between Earth and an orbiter might still utilize CCSDS Telemetry and Telecommand protocols, but between a rover and an orbiter it might make sense to use the new CCSDS Proximity-1 protocol [6] instead.

Elements of Interplanetary Internet

Operation of CFDP in these more complex scenarios is topologically similar to the operation of the Internet, except that the Internet's TCP and IP protocols as commonly implemented are defeated by large signal propagation delays and interruptions in end-to-end connectivity. In fact, the operation of CFDP over an Internet UT layer is subject to exactly the same constraints as basic deployment of CFDP in space, where the end-to-end route provided by IP functions as a single "link" from the CFDP point of view:

- The link must be topologically possible. That is, the subnets to which the CFDP endpoint entities are connected must be continuously connected (directly or indirectly) to the same backbone.

- The link must be operationally possible. That is, network configuration parameters at the CFDP endpoint entities must be correct.
- The link must be electronically possible. That is, network interface cards at the CFDP endpoint entities must be patched into active Ethernet ports.

Several years ago a study team began investigating the possibility of extending Internet capabilities into interplanetary space to construct an Interplanetary Internet (IPN) [7]. The approach taken by this study team has been to devise general-purpose delay-tolerant networking (DTN) technology that could function in all the configurations for which CFDP is currently being extended, i.e., transient violations of all of these constraints. The potential usefulness of DTN technology in fact seems to reach beyond the interplanetary domain: as connectivity is increasingly extended to novel, often stressed communication regimes on Earth, the ability to operate in terrestrial environments that are deep-space-like in some ways is becoming important for the Internet itself.

The DTN-based architecture contemplated for the IPN offers a superset of the features provided by either of the proposed approaches for extending CFDP. In addition to supporting store-and-forward transmission across a series of links, the DTN *bundling* protocol includes:

- Built-in facilities for data authentication and confidentiality.
- Routing algorithms that are far more flexible and dynamic than those planned for the CFDP extensions.

These features will be vital to the large-scale deployment of bundling in the Internet, and the same scalability will help enable the Interplanetary Internet to grow from a set of interoperating missions into a true network. Moreover, extensive use of DTN technology on Earth will tend to minimize the life cycle cost of deploying that technology in the IPN.

CFDP Advanced Deployment in Adapted Architecture

So the Interplanetary Internet offers a third approach for operating CFDP in complex scenarios: a UT layer interface to the DTN bundling protocol.

In an IPN environment, each CFDP PDU is encapsulated in a “bundle” (a bundling protocol datagram) and independently routed in store-and-forward fashion through the Interplanetary Internet. Bundle protocol agents at relay points forward the bundles; the files transmitted by CFDP are never delivered to CFDP entities – or CFDP user applications – at any relay point. CFDP is active only at the endpoint entities, just as in the basic deployment scenario, and CFDP point-to-point retransmission functionally devolves to end-to-end acknowledgment.

Parallel Relay Operations

In addition to enhanced security and more powerful routing, this DTN-adapted CFDP architecture offers straightforward support for operations in even more complex scenarios that involve transmission through multiple relay points in parallel.

This capability is highly desirable when, for example, CFDP at a Mars rover is required to send a very large file to Earth: it can’t transmit the whole file in a single relay orbiter tracking pass, but it can transmit part of the file in the next pass of one orbiter and the rest in the upcoming tracking pass of a second orbiter. CFDP has no built-in mechanism for distributing retransmission responsibility across multiple entities, though, so currently this scenario is not supported. The rover’s CFDP would have to wait for a second pass of the first orbiter in order to transmit the rest of the file.

A simple work-around is to divide any large file into multiple small files, such that each file can be sent in its entirety during a single tracking pass, and reassemble the large file at the final destination from the separately received small files. CCSDS is beginning to design a *data product management* standard user operation that would do just this.

However, a UT layer interface to DTN bundling would yield the same effect without any additional development: the independently routed bundles in which CFDP PDUs are encapsulated would be transmitted as tracking pass time becomes available, to whichever orbiter is currently overhead, and CFDP PDU content would be reassembled into the original file at the final destination in the usual way.

Status

A partial prototype implementation of the DTN bundling protocol has been developed. In the same way that CFDP runs over an abstract UT layer that can be mapped to various underlying protocols, the bundling protocol runs over an abstract *convergence layer*; convergence layer interfaces to TCP/IP and to a Sensor Network protocol have been developed and tested. The DTN prototype currently supports deferred transmission and store-and-forward operation using non-volatile storage. Implementations of DTN authentication, confidentiality, and schedule-sensitive routing are planned for fiscal 2003.

A UT layer interface to this prototype bundling implementation has been developed for JPL's implementation of CFDP, and CFDP running over bundling was informally demonstrated at JPL on May 23, 2002.

Outlook

CFDP is a stable international standard that can reduce mission operations cost and risk by enabling reliable file transfer and remote file system management over interplanetary distances. The application of emerging delay-tolerant networking technology to Interplanetary Internet operations, and specifically to the use of CFDP in complex mission configurations, will further enhance CFDP's usefulness and value to deep space exploration missions.

References

[1] *CCSDS File Delivery Protocol (CFDP)*. Recommendation for Space Data System Standards, CCSDS 727.0-B-1. Blue Book. Issue 1. Washington, D.C.: CCSDS, January 2002.

[2] <http://www.ccsds.org>

[3] *Packet Telemetry Service Specification*. Recommendation for Space Data System Standards, CCSDS 103.0-B-2. Blue Book. Issue 2. Washington, D.C.: CCSDS, June 2001.

[4] *Telecommand Part 2 – Data Routing Service*. Recommendation for Space Data System Standards, CCSDS 202.0-B-3. Blue Book. Issue 3. Washington, D.C.: CCSDS, June 2001.

[5] <http://deepimpact.jpl.nasa.gov/>

[6] *Proximity-1 Space Link Protocol*. CCSDS 211.0-R-1. Red Book. Issue 3. Washington, D.C.: CCSDS, January 2002

[7] <http://www.ipnsig.org>